

23

A No Fee

UTILITY PATENT APPLICATION

HONORABLE COMMISSIONER OF
PATENTS AND TRADEMARKS

Washington D.C. 20231

Case Docket No. 2860-065

Sun File No. P2287/TJC

Sir:

Transmitted herewith for filing is the patent application of

INVENTOR: Amit Gupta and Raphael Rom

FOR: TECHNIQUES FOR IMPROVING VIRTUAL CHANNEL MANAGEMENT AND
MAINTENANCE IN A NETWORK ENVIRONMENT

Enclosed are:

- ☒ 44 pages of specification, claims, abstract, and cover sheet
- ☐ Declaration & Power of Attorney.
- ☐ Priority Claimed
- ☒ 23 sheets of formal drawing.
- ☐ An assignment of the invention to _____
and the assignment recordation fee.
- ☐ Certified copy of _____
- ☐ An associate power of attorney.
- ☐ Information Disclosure Statement, Form PTO-1449 and reference.
- ☐

LOWE PRICE LEBLANC & BECKER

2860



David L. Stewart
Registration No. 37,578

99 Canal Center Plaza, Suite 300
Alexandria, Virginia 22314
(703) 684-1111 DLS:kap
Facsimile: (703) 684-1124
Date: June 4, 1997

06/04/97
71476 U.S. PTO

08868972-060497

Attorney Docket No. 2860-065; P2287/TJC

APPLICATION FOR UNITED STATES LETTERS PATENT FOR

TECHNIQUES FOR IMPROVING VIRTUAL CHANNEL MANAGEMENT
AND MAINTENANCE IN A NETWORK ENVIRONMENT

INVENTORS:
Amit Gupta
Raphael Rom

PREPARED BY:
David L. Stewart
LOWE PRICE LEBLANC & BECKER
99 Canal Center Plaza, Suite 300
Alexandria, VA 22314
Telephone: (703) 684-1111
Facsimile: (703) 684-1124

2860-065

P2287/TJC

**TECHNIQUES FOR IMPROVING VIRTUAL CHANNEL
MANAGEMENT AND MAINTENANCE IN A NETWORK ENVIRONMENT**

BACKGROUND OF THE INVENTION

Field of the Invention

The invention relates to communication systems and, more particularly to the creation, maintenance, use and deletion of virtual channels in a network environment.

5 Description of Related Art

Modern data networks trace their historical origins to telephone switching networks. Telephone switching networks set up connections between an originating caller and a called subscriber selected by dialing the called subscriber's telephone number.

Digital networks are known which permit information packets to be routed from a source to a destination without the necessity for a hardwired connection between the calling subscriber and the destination subscriber. Packets are routed over links between nodes of a network. Rather than a physical connection, a virtual connection,

sometimes called a virtual channel or virtual circuit is defined between the source of the packet and its destination. A virtual circuit gives the appearance of maintaining a hardwire connection, but utilizes the resources of the connection only when data need to be sent. This permits the link to be shared by other virtual circuits and improves the efficiency and throughput of the links of the network.

One form of modern high speed network uses a set of standards collectively termed asynchronous transfer mode (ATM). The basic data element used by these networks is a cell, which is a sequence of 53 bytes of which the first 5 are the header and the other 48 are the cell payload. To allow very high data rates (in the hundreds of megabytes per second and above) the switching time of a cell in every switch or node along the route through the network must be minimal. ATM networks achieve this by way of a virtual channel mechanism which requires a set-up phase before the actual data stream can be transported from source to destination. The notion of a virtual path is defined in the ATM standard. Each virtual path "contains" 2^{16} (65,536) virtual circuits. Supporting virtual path requires significant additions to the ATM switch hardware. Consequently, many switches only support virtual circuits; they do not support virtual paths.

The Problems

It would be desirable to have an efficient mechanism for quickly allocating multiple connections without requiring extra hardware support in the network nodes.

5 Further, it would be desirable to reduce the table size required for network switches to support large numbers of virtual circuits. It would further be desirable to implement multiple virtual channels set-up while maintaining compatibility with current switches but

10 nevertheless supporting heterogeneity in switches in a network. It would further be desirable to be able to route multiple virtual circuits over different paths. It would also be desirable to support dynamic aggregation of individual virtual circuits into larger entities. It

15 would also be desired to permit enhanced resource sharing across connections of a network. It would further be desirable to have techniques for handling ATM cell interleaving problems in a multi-sender multi-cast. It would also be desirable to have efficient support for

20 fast connection establishment. It would further be desirable to improve the efficiency with which "keep-alive" and "refresh" packets are sent through the network.

SUMMARY OF THE INVENTION

The problems of the prior art are solved in accordance with the invention by providing a new entity for use in networks called the Virtual Circuit Bunch (VCB). Virtual Circuit Bunches can be utilized to set-up and manage together groups of virtual circuits in a flexible way which results in increased performance as well as overcoming the problems of the prior art.

The Virtual Circuit Bunch, as set forth herein, enables groups of Virtual Circuits to be established between one or more points of origin and one or more destinations. Virtual Circuit Bunches can be implemented without any changes to switch hardware and permit great flexibility in Virtual Circuit allocation.

One aspect of the invention is directed to a switching node which includes a switching matrix, and a controller to control the switching matrix, in which the controller is configured to set up group(s) of virtual circuits to one or more destinations as a virtual circuit bunch.

Another aspect of the invention is directed to a method of allocating virtual circuits, by establishing a plurality of virtual circuits from one node to at least one other node as a virtual circuit bunch in response to a single request.

Another aspect of the invention relates to a method of identifying virtual circuits at a node going to a common destination node; and aggregating those virtual circuits into a virtual circuit bunch.

5 Another aspect of the invention relates to a method of providing a fast connect service in a digital switching network by appropriately assigning a packet to a virtual circuit of a virtual circuit bunch without establishing a connection.

10 Another aspect of the invention relates to a method of allocating virtual circuits in a switching system allocating a virtual circuit to all nodes participating in a multicast using a virtual circuit bunch.

15 Another aspect of the invention relates to a system for the transmission of digital communications, including a plurality of user communication devices; a plurality of at least partially interconnected switching nodes, each node serviced by a node controller servicing said user communications devices; in which at least one of said
20 node controllers is configured to set up a group of virtual circuits to respective one or more destinations as a virtual circuit bunch.

Other aspects of the invention are directed to computer program products for carrying out the invention.

25 The foregoing and other features, aspects and advantages of the present invention will become more

apparent from the following detailed description of the present invention when taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The objects, features and advantages of the present invention will be apparent from the following description, in which:

Figure 1 has a block diagram showing an exemplary network used for carrying out the invention.

10 **Figure 2** is a block diagram of an exemplary node of the network shown in **Figure 1**.

Figure 3A is a front view of a computer suitable for use as a control device or controller for one or more network nodes.

15 **Figure 3B** is a block diagram showing an exemplary hardware architecture of the controller of **Figure 3A**.

Figure 3C is a view of an exemplary memory medium suitable for storing programs and data used in carrying out the invention.

20 **Figure 4** is a flow chart of a process used at a node for setting up a virtual circuit from an originating station U1 to a destination station U2.

Figure 5 is a time line diagram showing set-up of a virtual circuit between U1 and U2 over the network of **Figure 1**.

25

Figure 6 is a block diagram showing operation of a network node when implemented as a ATM switch.

Figures 7A, 7B and 7C show exemplary packet formats used in requesting set-up of a Virtual Circuit Bunch (VCB).

Figure 8A is a table summarizing the VCB request of **Figure 7C** and **Figure 8B** shows the same table but with the incoming ports utilized in the VCB assigned to carry traffic.

Figure 9 is a block diagram showing the switches and links used for setting up the VCB requested in **Figure 7C**.

Figure 10 is an exemplary table storing VCB information.

Figure 11 shows exemplary tables storing connection information relating to a VCB used at a node at either end of a VCB.

Figure 12 shows exemplary tables storing connection information relating to a VCB at a node intermediate to the end nodes of a VCB.

Figure 13 is a flow chart of a process for setting up a VCB in accordance with the invention.

Figure 14 is a flow chart of a process for setting up a VC over an established VCB in accordance with the invention.

Figure 15 is a flow chart of a process for breaking down a VCB in accordance with the invention.

Figure 16 is a flow chart of a process for aggregating virtual circuits into a VCB in accordance with the invention.

Figure 17 is a block diagram illustrating statistically sharing virtual circuits of a VCB for a fast connect service.

Figure 18 is a flow chart of a process used for providing a fast connect service.

Figure 19 is a block diagram of a portion of a network used to illustrate split routing of a VCB.

Figure 20 is a block diagram showing use of a VCB when multi-casting over a network.

Figure 21 is a block diagram showing use of a VCB when multi-casting using plural sources over a network.

Figure 22 is a block diagram showing the cell interleaving problem and how it is solved when using a VCB.

NOTATIONS AND NOMENCLATURE

The detailed descriptions which follow may be presented in terms of program procedures executed on a computer or network of computers. These procedural descriptions and representations are the means used by those skilled in the art to most effectively convey the substance of their work to others skilled in the art.

A procedure is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. These steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It proves convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be noted, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

Further, the manipulations performed are often referred to in terms, such as adding or comparing, which are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary, or desirable in most cases, in any of the operations described herein which form part of the present invention; the operations are machine operations. Useful machines for performing the operation of the present invention include general purpose digital computers or similar devices.

The present invention also relates to apparatus for performing these operations. This apparatus may be

5 specially constructed for the required purpose or it may
comprise a general purpose computer as selectively
activated or reconfigured by a computer program stored in
the computer. The procedures presented herein are not
inherently related to a particular computer or other
apparatus. Various general purpose machines may be used
with programs written in accordance with the teachings
herein, or it may prove more convenient to construct more
specialized apparatus to perform the required method
10 steps. The required structure for a variety of these
machines will appear from the description given.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

15 **Figure 1** is a block diagram showing an exemplary
network useful in carrying out the invention. The
network shown is comprised of nodes A through J, each
node being capable of routing digital data from users
connected to the node or from other nodes connected to
the node through to a destination. For example, when
user U1, which is connected to node A, desires to connect
20 to user U2, connected to node J, the user will request a
connection from node A to node J as discussed more
hereinafter. In one implementation, node A will do a
calculation of a proposed route through the network from
user U1 to user U2. In the example shown, node A
25 proposes a route from user 1 through node A, to node C,

to node F, to node H, to node J, to user U2. Different types of networks have different types of routing mechanisms. In the example shown in **Figure 1**, the invention will be described with reference to an ATM network where each of the nodes constitutes an ATM switch. The principles of the invention, however, are applicable to general networks without restriction to the type of switching technology.

Figure 2 is a block diagram of an exemplary node of the network shown in **Figure 1**. Digital data arrives on incoming ports 1 through N over digital links connected to users and other nodes. The incoming digital data is optionally buffered in buffers **240** prior to being routed to the appropriate outgoing port of the set of outgoing ports 1-M. The outgoing ports connect to either end-users or other nodes as desired to implement a particular network topology. One or more interconnect switches **230** forming part of a switching fabric are activated to connect the appropriate incoming port to the appropriate outgoing port so that data is routed from the source toward the destination in an orderly fashion. Data having passed through the switching fabric **230** is optionally applied to outgoing buffers **250** prior to being applied to respective links connected to the outgoing ports. A control device or controller **200**.

At a high level, ATM networks can be viewed as a collection of ATM-switches interconnected by high-speed links to provide ATM service to a number of external users. A typical switch is comprised of three main parts: the switching fabric, a buffering unit, and a control unit. The control unit is divided into two parts. The first part controls the cell switching properties, i.e., controls and coordinates the operation of the buffering unit and the switching fabric; for performance reasons it is typically implemented in hardware such as lookup tables, finite state machines, etc. The second part is responsible for all other control functions such as preparing and maintaining the switching tables, coordinating activities with other switches and with network managements facilities, and servicing the individual user requests. This is referred to as the control point (CP) of the switch. As used herein the switch contains the fabric, the buffering units, and the basic core of the control unit, whereas the switch and the CP together will be referred to as a node. The CP is essentially a logical entity which is typically implemented in software although firmware and other software/hardware configuration are also possible.

The device on which the CP is implemented is teamed the Controlling Device (CD) or controller. The CD may not be a single entity, that is, parts of the CP can be

Every switch in the network possesses a unique ID, which is assigned to it at configuration time. This ID, termed switch ID, is guaranteed to be unique only within the network. Links can have IDs that are local to the switch from which they emanate. This allows to keep the local link IDs (which are frequently used) short. Therefore, to globally identify a link some combination of a CP ID, node ID and link ID must be used. Link, switch and CP IDs appear in the messages that are exchanged between CPs and hence it is desirable to choose as short a descriptor as possible.

In general, CPs perform those functions that are necessary to make a group of switches operate as an integrated network. The process of integration entails understanding of and compensation for the differences between switches of different makes and models. One of the most important functions is that of virtual circuit (connection) management which includes VC setup, maintenance and VC take down.

To perform their functions, CPs obviously need to communicate with one another either directly or indirectly. The control mechanism typically requires CPs to mostly communicate with neighbor CPs. In many cases a message arriving at a CP may trigger the transmission of another message. For example, if a topological change has occurred, a CP will typically notify other CPs of the

change; each of the recipients of this message may, in turn, generate new messages to further distribute the information. A control link as the channel over which two CPs communicate. The collection of control of control links must enable every CP to distribute the appropriate data to every other CP. In a typical design, there will be a control link between any two neighboring CPs.

Figure 3A is a front view of a computer suitable for use as a control device or controller for one or more network nodes. Although the computer is illustrated as a standalone device, computers used as controllers are more typically printed circuit cards which are rack mounted and plug into a backplane. The controller **300** illustrated includes drives **310A** and **310B** for receiving, reading and writing information stored on removable storage media. The controller **300** also may optionally include a CD ROM device, not shown and an internal hard drive, also not shown. Display **320**, keyboard **330** and mouse **340** may be optional in a particular environment although typically, these input output devices may be attached for troubleshooting a particular controller. The actual configuration of a controller may vary from switch to switch.

Figure 3B is a block diagram showing an exemplary hardware architecture of the controller of **Figure 3A**.

As shown in **Figure 3B**, the controller has a CPU 355 which connects and controls bus 350. A disk controller 370 interfaces the floppy drives 373, internal hard drive 372 and CD ROM 371. The particular configuration for storage media may vary from controller to controller. Read Only Memory 360 and Random Access Memory 365 provide memory for storing programs and data information to carry out the controller functions. Interfaces 345 and 375 are available to permit input output devices such as keyboard 330, mouse 340 and display 320 to be connected when working with the controller for installation, setup, maintenance or the like. A communications port 385 permits external communications from the bus. Switch interface 190 provides the control and exchange of information between the switch controlled by the controller and the CPU 355 providing the control.

Figure 3C is a view of an exemplary memory medium suitable for storing programs and data use in carrying out the invention.

Figure 4 is a flow chart of a process used at a node in setting up a virtual circuit from an originating station U1 to a destination station U2. This process describes a simple single VC connection between U1 and U2. When a connection request is received on an incoming port i, such as from a user or a predecessor node, it is routed to the controller as discussed previously (400).

The controller consults a network map and determines the (best) next node for routing data from U1 toward its destination address (410). In some networks, the controller will only determine the best next node whereas in others, it will attempt to determine the entire route to the destination. The controller selects an outgoing port and VCI to the next node with adequate resources to service a level of service (LOS) specified for the connection (420). The controller sends a request control packet on the outgoing port and virtual circuit to the next node controller and waits for an acknowledgment (430). The controller may send this over a control link between the two switches or over the actual port and virtual circuit proposed for use in the connection. When an acknowledgment is received from the next node, the acknowledgment will be sent to predecessor node or to the user indicating acceptance of the connection and a willingness to provide the services at the requested LOS specified in the request. The node controller will make entries in the switch table specifying the port and VC to be utilized for the connection (440).

Figure 5 is a time line diagram showing set-up of a virtual circuit between U1 and U2 over the network of Figure 1. User U1 initiates a connection request message to node A (500) which requests a virtual circuit to node C (505) which requests a virtual circuit to node F (510)

which requests a VC to node H (515) which requests a VC to node J (520) which requests a VC to destination user U2 (525). Once the VC is dedicated to node U2, U2 will acknowledge the VC to node J (530) which will acknowledge the VC to node H (535) which will acknowledge the VC to node F (540) which will acknowledge the VC to node C (545) which will acknowledge the VC to node A (550) which will acknowledge the connection established (555) to the requesting user U1. Thus a cascading set of VC requests normally occurs followed by a cascading set of acknowledgments by which a virtual circuit is set up. One should note that this requires significant processing at each of the nodes and significant network traffic in order to set up a single VC. If one is attempting to set up a thousand virtual circuits, individual handling of the virtual circuits creates a significant processing and network traffic load.

Figure 6 is a block diagram showing an operation of a network node implemented as an ATM switch. In order to quickly switch an ATM cell arriving at an incoming buffer, such as 640 to an outgoing port, a set of switching tables is maintained. Each switching table has an (incoming) virtual circuit identifier (VCI) column and an associated output port (and VCI). Virtual Circuit Identifier's are sometimes called tags. In **Figure 6**, the ports illustrated are bi-directional ports, that is, they

accommodate incoming as well as outgoing traffic. When an incoming ATM cell is received in an in-buffer for port 3, the tag (XXXXXX) is read and utilized to extract the output port destination for the tagged cell using the port 3 forward table 660. Port numbers passing to the switching interface 600 activate the switching fabric 630 to connect port 3 to port 22 during which time the contents of in-buffer 640 are passed to an out buffer 650 on port 22. In the reverse direction, the tag of an incoming cell stored in in-buffer 650 (YYYYY) is utilized to access the port 22 reverse table 670. The input port, 3 in this case, is passed to a switching interface 600 to set up a reverse direction connection between the port 22 in-buffer and port 3 out-buffer 640. The switch tables are set up for a very fast access and response. The contents of the switching tables as illustrated in Figure 6 are set by the controller from the controller tables as discussed more hereinafter. In this example, the bi-directional VC takes the same route in both directions. This does not need to be the case, but is for purpose of this example.

Figures 7A, 7B and 7C show exemplary packet formats used in requesting set up of a virtual circuit bunch (VCB). The generic format for such a request is shown in Figure 7A. A packet header 700 begins a request with a packet type 710 indicating a VCB request. Following the

VCB request are a set of ordered fields **720A1**, **720B1** and **720C1** through **730AN**, **730BN** and **730CN**. Each set of the ordered set of fields specifies a destination (**720A**), a number of VCs (**720B**) for a virtual circuit bunch or subset of virtual circuit bunch going to the destination specified in **720A**, and a level of service (LOS) **720C**. A single virtual circuit bunch request can be utilized to specify a plurality of destinations and the number of VCs associated with each destination can be arbitrarily specified.

Figure 7B shows an example of a request for a virtual circuit request from switch 1 to switch 2 for 32 virtual circuits. This would be sent as a control packet from the switch 1 controller to the switch 2 controller.

In more complex example, one which will be pursued throughout the remainder of the explanation of the invention, is shown in **Figure 7C**. Again a packet header **700** and a VCB request field **710** precede the fields specifying the virtual circuit bunch. In the example shown in **Figure 7C**, switch 1 is requesting a virtual circuit bunch totalling 25 virtual circuits, 8 of which are destined for switch 4, 8 of which are destined for switch 5, and 9 of which are destined for switch 3. When the control packet illustrated in **Figure 7C** is received at switch 2, switch 2 acts strictly as an intermediary to traffic going to switches 3, 4, and 5. This does not

need to be the case, but is for purposes of this example. Switch 2 could easily be a destination for some of the virtual circuits. By this request, switch 1 sets up 25 bi-directional virtual circuits to switch 2. Switch 2 sets up 8 outgoing virtual circuits to switch 4, 8 to switch 5 and 9 to switch 3. In a reverse direction, the virtual circuits from switches 3, 4, and 5 are mapped to the 25 virtual circuits set up between switch 1 and switch 2. When the virtual circuit bunch is set up, no actual traffic is passed, rather the virtual circuits are defined together with their level of service and are available to pass traffic as it is routed to them from circuits originating or terminating at the end nodes of the virtual circuit bunch.

Figure 8A is a table summarizing the VCB request of **Figure 7C**. Considering columns 3 and 4 of **Figure 8A**, a virtual circuit bunch of 25 virtual circuits is set up between switch 1 and switch 2. This virtual circuit bunch is labeled VCB 3 for reference. VCB 3 is a convenient alias for referring to groups of virtual circuits, in this case consisting of 25 virtual circuits between switch 1 and switch 2. At switch 2, portions of the virtual circuit bunch VCB 3 are directed to different locations. Specifically, 8 VCs are directed to switch 4, 8 are directed to switch 5, and 9 are directed to switch 3, these subsets of VCB 3 can be labeled, conveniently,

VCB 3A, VCB 3B and VCB 3C. These labels are convenient aliases for the subset(s) of virtual circuits going respectively to destinations switch 4, switch 5, and switch 3, respectively.

5 **Figure 8B** shows the same table as **Figure 8A** except that exemplary port and VC assignments are shown.

Figure 9 is a block diagram showing the switches and links used in setting up the VCB requested in **Figure 7C** and described in conjunction with **Figure 8**. VCB 3 is
10 established between switch 1 and switch 2 with 25 virtual circuits. Those 25 virtual circuits are divided up into 3 subsets with subset VCB 3A going to switch 4, subset VCB 3B going to switch 5, and subset VCB 3C going to switch 3.

15 At this point in the example, the request of **Figure 7C** has been granted and the virtual circuit bunch established as illustrated in **Figure 9**. However, no traffic is flowing across the virtual circuit bunch. The virtual circuit bunches are allocated and set up in the
20 controller tables and the switching tables of the switches but they are not being utilized.

 Returning to **Figure 8B**, column 1 of **Figure 8B** shows an allocation of incoming ports and virtual circuits of switch 1 which correspond to the virtual circuit bunch
25 VCB3 and its subsets VCB 3A and VCB 3B and VCB 3C. When a new user at switch 1 has traffic for a destination at

switch 4, switch 1 will assign the user to a VC in VCB 3A which goes to switch 4. The assigned VC in VCB 3A will immediately route the traffic to switch 4 without any setup required for virtual circuits between switch 1 and switch 4. They have been preallocated. At switch 4, the destination address of a cell is used to route the cell to its final destination.

Figure 10 is an exemplary controller table storing VCB information. This one is located at switch 1. At switch 1, VCB 3 is defined as a set of <port, VCI> tuples listed in the right hand column of the table together with the destination for those <port, VCI> tuples. The entire set of 25 VCBs may be referred to by the convenient alias VCB 3. Alternatively, subsets of VCB 3 can be referred to by the convenient aliases VCB 3A, VCB 3B, and VCB 3C, each having a respective destination. Thus, a hierarchy of aliases, which may number more than the two levels shown, maybe utilized to conveniently refer to portions of a group a virtual circuits preestablished in a virtual circuit bunch.

Figure 11 shows exemplary tables in a switch controller storing connection information relating to a VCB used at a node at either end of a virtual circuit bunch. Different tables are utilized at intermediate nodes. These tables are utilized by the switch controller to establish switching table values for

controlling the switching fabric to route an incoming port and VCI to its appropriate outgoing port and VCI. As shown in **Figure 11**, different tables are utilized for the forward and reverse direction. Referring to the table shown in **Figure 8B**, table 1100 which is for the forward direction differs from table 1120 for the reverse direction. Each of these table entries only shows the <port, VCI> tuples associated with the virtual circuit bunch although other entries are present. As shown in column 1 of **Figure 8B**, virtual circuits 102-106 and 111-113 from port 3 have requested connections to destination node switch 4. Thus they are assigned a virtual circuit identifier, in column 1140 of **Figure 11**, VCB 3A. The specific output <port, VCI> tuples to be utilized for these virtual circuits is any unused VCI on outgoing port 5 in the VCB 3A. Similarly, the other lists of input <port, VCI> tuples and virtual circuits found in column 1130 of table 1100 have an alias identifier of the virtual circuit bunch to which they have been assigned for routing data to the destination serviced by the respective virtual circuit bunch.

The reverse direction is specified in table 1120. Since any <port, VCI> tuple associated with VCB 3 may route to any destination on switch 1 the destination address for the cell is read and the corresponding <port, VCI> specified for the destination is utilized to control

the switching fabric so that data incoming on VCB 3 is routed to the correct destination.

Considering these tables in more detail, when table 1100 is utilized and the VCI is listed as a virtual circuit bunch alias, the definition of the virtual circuit bunch is looked up in the table described in **Figure 10**, selection of an unused <port, VCI> tuple in the virtual circuit bunch is made and an unused <port, VCI> tuple is written into the forward and reverse switch tables to set up a virtual circuit connection.

Figure 12 shows exemplary tables storing connection information relating to a VCB at a node intermediate to the end nodes of a VCB. This table is one typically found in an intermediate switch, such as switch 2 shown in **Figure 9**. Table 1200 shown in **Figure 12** is very similar table 1100 shown in **Figure 11**.

Note that table 1200 shown in **Figure 12** contains many fewer entries than would be required if each virtual circuit were set up separately without the benefit of virtual circuit bunches. Table 1210 shows an even greater conciseness of expression. A single entry is sufficient to specify the reverse direction for all 25 virtual circuits in the virtual circuit bunch. In the reverse table, data incoming on any of the listed virtual circuits would be routed to VCB 3 and will be assigned to any unused virtual circuit in VCB 3 for transmission to

switch 1. Again to determine the unused output <port, VCI> tuple in VCB 3, one could refer to the table, such is shown in Figure 10 at switch 2 and then check occupancy data for the virtual circuits participating in VCB 3.

Figure 13 is flow chart of a process for setting up a VCB in accordance with the invention. The process begins when a user or node requests a VCB setup with a control packet to the local node control point (1300). The local node control point parses the control packet and calculates routes to the destination(s) specified (1310). This may be either a complete route calculation or a best next node calculation. The controller identifies outgoing virtual circuits needed to satisfy the contract for the level of service needed for the connections requests to each destination (1320). The local node control point then requests allocation of a virtual circuit bunch to the next node along with path, actually specifying the VCs desired (1330).

When an acknowledgment is received from the next node, an acknowledgment is sent either to the predecessor node or the requesting station 1340 and table entries are created or finalized (1350) thus setting up the virtual circuit.

Figure 14 is a flow chart of a process for setting up a virtual circuit using an established virtual circuit

bunch in accordance with the invention. Continuing with the previous example, the switch 1 controller receives a request for a VC to a destination switch N serviced by the VCB (1400). The switch 1 controller sends a control packet to the switch N controller either over a VC of the VCB or over a control link (1410) requesting a connection to a destination address at switch N. The switch N controller enters controller table data for the <port, VCI> tuple to be used by the destination address, sets the switch table and acknowledges the connection by sending a control packet to the switch 1 controller (1420). The switch 1 controller enters the controller table data for the destination address and the originating address and sets the switch table entries (1440), thus establishing an end to end connection between user 1 and user 2. Thus, in contrast to the process described in conjunction with Figures 4 and 5, since a virtual circuit bunch has already been set up, the only action required to establish an end to end connection is at the end points of the virtual circuit bunch. Thus, the switch 1 controller needs only establish a connection between the requesting virtual circuit for user U1 and the virtual circuit bunch destined for switch N, and switch N needs only to establish an association between the virtual circuit bunch and the destination user U2. Thus network traffic

is significantly reduced, processing at the intermediate nodes is eliminated since the virtual circuit bunch has already been set up.

Figure 15 is a flow chart of a process for breaking down a virtual circuit bunch in accordance with the invention. A virtual circuit bunch breakdown request is received on an incoming port *i* and routed to the controller **(1500)**. The controller consults the table for entries using aliases of the VCB being broken down and marks them for deletion **(1510)**. The connection breakdown request is forward to the next node, if any; otherwise table entries for the VCB being broken down are deleted and acknowledgment is sent to the requesting node **(1520)**. When acknowledgment of a breakdown requests is received, table entries have been actually deleted and all traces of the virtual circuit bunch are eliminated from the system **(1530)**.

Figure 17 is a block diagram illustrating sharing of virtual circuits of a virtual circuit bunch for a fast connect service. For purposes of this illustration assume that users U1A through U100A generate messages of short duration. A five circuit bunch VCB **1700** is shown going between the node servicing users U1A through U100A and the node servicing users U1B through U100B. Block **1710** shows a distribution mechanism which will assign the short duration message from users U1A through U100A to

the five virtual circuits of the VC bunch **1700** according to a distribution policy that may be partially deterministic, random, or pseudo-random. In any case, at the destination end, **1720**, the destination addresses of the cells are read and the cells routed to the destination address without a formal virtual circuit being setup. In short, when one of users U1A through U100A has a message to send, it is assigned to one of the virtual circuits of virtual circuit bunch according to the distribution policy without actually setting up a connection and it is routed at the receiving end to the destination, also without setting up a connection. Thus, the five virtual circuits of the virtual circuit bunch can service a much larger number of users without a separate on the fly connection establishment. This provides a very quick and efficient fast connect service.

Figure 18 is a flow chart of the process described in conjunction with **Figure 17**. When a fast connect packet plus request is received at a node, a <port, VCI> tuple from a VCB going to the destination node is selected (**1700**). The fast connect packet/request is sent over the selected <port, VCI> tuple to the destination node (**1710**). The fast connect packet/request is routed at the destination node to the specified address (**1720**) and an optional fast connect acknowledgment is returned

to the originating address in the same way fast connect packet/request was sent (1730).

Figure 19 is a block diagram of a portion of a network used to illustrate split routing of a VCB. Assume that node A requests 150 virtual circuits to destination E at a given level of service (LOS). Only 100 such virtual circuits are available through C and node D has only 85 available. This decision may also be the result of a policy decision (for example, load-balancing) of the network control system. Node B can allocate a virtual circuit bunch routing a portion of the VCB over a node C and the remainder over node D. This kind of split routing was not a capability provided by the virtual path technology of the prior art. Thus VCBs can split the routing of a virtual circuit bunch into convenient units.

The same figure illustrates another benefit achieved by the invention. In the prior art, there is sometimes a need to send keep alive packets or refresh packets periodically over each virtual circuit to avoid a connection timing out and being broken down. With a virtual circuit bunch, one only needs to send one keep alive packet or refresh packet per virtual circuit bunch and not one for each circuit. This reduces network traffic and processing power requirement.

Figure 20 is a block diagram showing the cell interleaving problem and how it is solved when using a virtual circuit bunch. The problem arises when, as shown in **Figure 20**, two <port, VCI> tuples attempt to send to the same destination. Thus, the port 1 switch table attempts to route VCI 120 to port 2 VCI 150. Similarly, as shown in table **2020**, port 4 attempts to route VCI 125 to outgoing port 2 and VCI 150. Thus, cells from two sources would occur on the same outgoing port and VCI and would result in a error condition at the destination. Such a failure will be detected when the CRC result will not conform to the expected result. If outgoing ports 2 and 6 were part of the same virtual circuit bunch, one of the two conflicting connections could be routed over another unused VCI of the virtual circuit bunch or over an unused VCI of port 2. Alternatively, some of the virtual circuits could be used to permit sharing by appropriate selection as discussed above. Thus, when cell interleaving occurs, rerouting can be easily accomplished using a virtual circuit bunch and cells successfully routed to the correct destination without error.

The previous examples illustrated VC bunching is somewhat static manner - the bunches were established as a unit, and they remained the same through the life-time. However, VC bunches can also be established and modified

dynamically. The system may decide that a set of already established VCs (and bunches) can be served more efficiently by aggregating them together into a large bunch. Alternately, when a new VC (or bunch) establishment request arrives, the switch may decide to incorporate it into an existing bunch (by expanding it). Such dynamic aggregation may require changes in VC identifiers.

Figure 21A illustrates this concept; initially, the switches A, C, D, E and F carry a bunch of size 32. Now, node B wants to set up a bunch of size 32 to node G. When the new bunch request arrives at node C (from node B), node C aligns the new VC bunch with the existing bunch, to create a size 64 bunch at entry to node D as shown in **Figure 21B** which sets up the expanded bunch of 64 to node E, where the two bunches are separated once again (as the paths diverge). In this manner, one can dynamically create flexible bunches via aggregation.

The invention disclosed provides improved efficiency for setting up multiple connections in a network environment. It results in a reduction in the virtual circuit table size. It maintains compatibility with current switches without requiring hardware changes to participate in the VC bunching. It permits use in multi-path virtual circuit bunches which was not feasible in the case of virtual paths. It supports aggregation by

which data paths may start out as individual circuits but later the neighbors may agree to change identifiers so as utilize VC bunching. It provides efficient support for sharing resource allocations. It supports a variety of techniques for handling ATM cell interleaving problems, particularly in multi-sender multi-cast. It provides efficient support for fast connection set up and for a fast connect service. It provides considerable efficiency for handling keep-alive and refresh packets. It does not require a full 2^{16} address space required by virtual paths. Only the address space required for a particular virtual circuit bunch is set up. It allows virtual circuit bunches to be subdivided to service multiple destinations. Once a virtual circuit bunch is set up essentially zero latency is required for set up of a virtual circuit from the origin to the destination. With virtual circuit paths, once the large address space was allocated, resources were assigned and dedicated to it. Those resources were not used, the resources were wasted. Virtual circuit bunches permit allocation of resources on an as needed basis.

Thus, the techniques of the invention described overcome the problems of the prior art and permit great flexibility in the establishment and use of virtual circuits.

5

What is claimed is:

1. A switching node, comprising:
 - a. a switching matrix, and
 - b. a controller to control said switching matrix,
said controller configured to set up at least one group
5 of virtual circuits to respective one or more
destinations as a virtual circuit bunch.
2. The switching node of claim 1 in which said
switching node is an ATM switch.
3. The switching node of claim 1 in which said
controller sets up a virtual circuit bunch by sending a
message to another controller specifying one or more
destinations and a respective number of virtual circuits
5 to go to each destination.
4. The switching node of claim 1 in which said
controller keeps all virtual circuits of a virtual
circuit bunch to a single destination alive by sending or
receiving a refresh packet on less than all of the
5 virtual circuits going to said single destination.

5. The switching node of claim 4 in which a single refresh packet is used to keep all virtual circuits of a virtual circuit bunch to a single destination alive.

6. The switching node of claim 1 in which said controller is configured to assign digital information from a source to one of a plurality of virtual circuits of a virtual circuit bunch.

7. The switching node of claim 6 in which the assignment of digital information from a source to one of a plurality of virtual circuits of a virtual circuit bunch is done without assigning said one of a plurality of virtual circuits to a connection.

8. The switching node of claim 1 in which the virtual circuits of a virtual circuit bunch going to a single destination may be routed over different paths.

9. The switching node of claim 1 in which the controller is configured to retransmit digital data from an assigned virtual circuit identifier (VCI) to an alternate VCI of the same or different port going to the same destination when a cell interleaving problem occurs.

10. Computer apparatus for connection to a switching node comprising:

- a. a bus;
- b. an input device, connected to said bus;
- 5 c. a communications interface connected to said bus;
- d. a processor, connected to said bus, said processor configured to receive an input from a user over said input device and to generate a single request to said switching node to establish a plurality of virtual circuits to respective one or more destinations as a
- 10 virtual circuit bunch.

11. In a digital switching network having a plurality of interconnected nodes, a method of allocating virtual circuits, comprising the step of:

- a. providing an element for performing the step of
- 5 establishing a plurality of virtual circuits from one node to at least one other node as a virtual circuit bunch in response to a single request.

12. The method of claim 11, in which the step of establishing a plurality of virtual circuits from one node to at least one other node as a virtual circuit bunch includes setting up switching tables when at least

5 one subsequent node has acknowledged the request.

13. The method of claim 11, in which said request specifies a plurality of destinations.

14. The method of claim 13, in which said request also specifies the number of virtual circuits to be established to each destination.

15. The method of claim 11, in which the request specifies the level of service to be provided by one or more virtual circuits.

16. The method of claim 11, further comprising the step of establishing an end to end virtual circuit using at least one virtual circuit of said virtual circuit bunch.

17. A method of allocating virtual circuits in a switching system, comprising the steps of:

a. identifying virtual circuits at a node going to a common destination node; and

b. aggregating those virtual circuits into a virtual circuit bunch.

18. A method of providing a fast connect service in a digital switching network, comprising the step of:

a. assigning a packet to a virtual circuit of a virtual circuit bunch.

19. The method of claim 18 in which a packet is assigned a VCI without setting up a connection.

20. The switching node of claim 1 in which said controller is configured to assign digital information from a source to one of a plurality of virtual circuits of a virtual circuit bunch in accordance with a user specified policy.

21. A system for the transmission of digital communications, comprising:

a. a plurality of user communication devices;

b. a plurality of at least partially interconnected switching nodes, each node serviced by a node controller, servicing said user communications devices;

c. in which at least one of said node controllers is configured to set up a group of virtual circuits to respective one or more destinations as a virtual circuit bunch.

22. The system of claim 21 in which a virtual circuit from a user at one node is connected to a user at a destination node using a virtual circuit from said virtual circuit bunch.

23. A computer program product, comprising:

a. a memory medium, and

b. a computer program stored on said memory medium, said computer program comprising instructions for establishing a plurality of virtual circuits from one node to at least one other node of a switching network as a virtual circuit bunch.

24. A computer program product, comprising:

a. a memory medium, and

b. a computer program stored on said memory medium, said computer program comprising instructions for identifying virtual circuits at a node going to a common destination node; and for aggregating those virtual circuits into a virtual circuit bunch.

25. A computer program product, comprising:

a. a memory medium, and

b. a computer program stored on said memory medium,
said computer program comprising instructions for
5 assigning a packet to a virtual circuit of a virtual
circuit bunch.

26. A computer program product, comprising:

a. a memory medium, and

b. a computer program stored on said memory medium,
said computer program comprising instructions for
5 allocating a virtual circuit to all nodes participating
in a multicast using a virtual circuit bunch.

27. A method of transferring the computer program
product of claim 23, comprising the step of transmitting
said instructions from said memory medium to a
destination over a communications interface.

28. A method of transferring the computer program
product of claim 24, comprising the step of transmitting
said instructions from said memory medium to a
destination over a communications interface.

29. A method of transferring the computer program product of claim 25, comprising the step of transmitting said instructions from said memory medium to a destination over a communications interface.

30. A method of transferring the computer program product of claim 26, comprising the step of transmitting said instructions from said memory medium to a destination over a communications interface.

20080607 10:04:37

**TECHNIQUES FOR IMPROVING VIRTUAL CHANNEL
MANAGEMENT AND MAINTENANCE IN A NETWORK ENVIRONMENT**

Abstract of the Disclosure

The invention is directed to a flexible method of allocating groups of virtual circuits in a network as a virtual circuit bunch. A virtual circuit bunch of arbitrary number can be set up with the ease of setting up a single virtual circuit. A single virtual circuit bunch can have plural destinations. Further, a virtual circuit bunch can be routed to the same destination over different routes. Connections across the network can be initiated as separate virtual circuits and then be grouped together as a virtual circuit bunch. The implementation of virtual circuit bunches permits rapid setup of virtual circuits and reduced table sizes at individual nodes of the network. A fast connect service can be implemented using virtual circuit bunches without establishing a connection by assigning a fast connect packet to one virtual circuit of a virtual circuit bunch. Virtual circuit bunches permit cell interleaving problems found in ATM switches to be circumvented in a particularly convenient way.




Figure 1




Figure 2





Figure 3A






Figure 3C

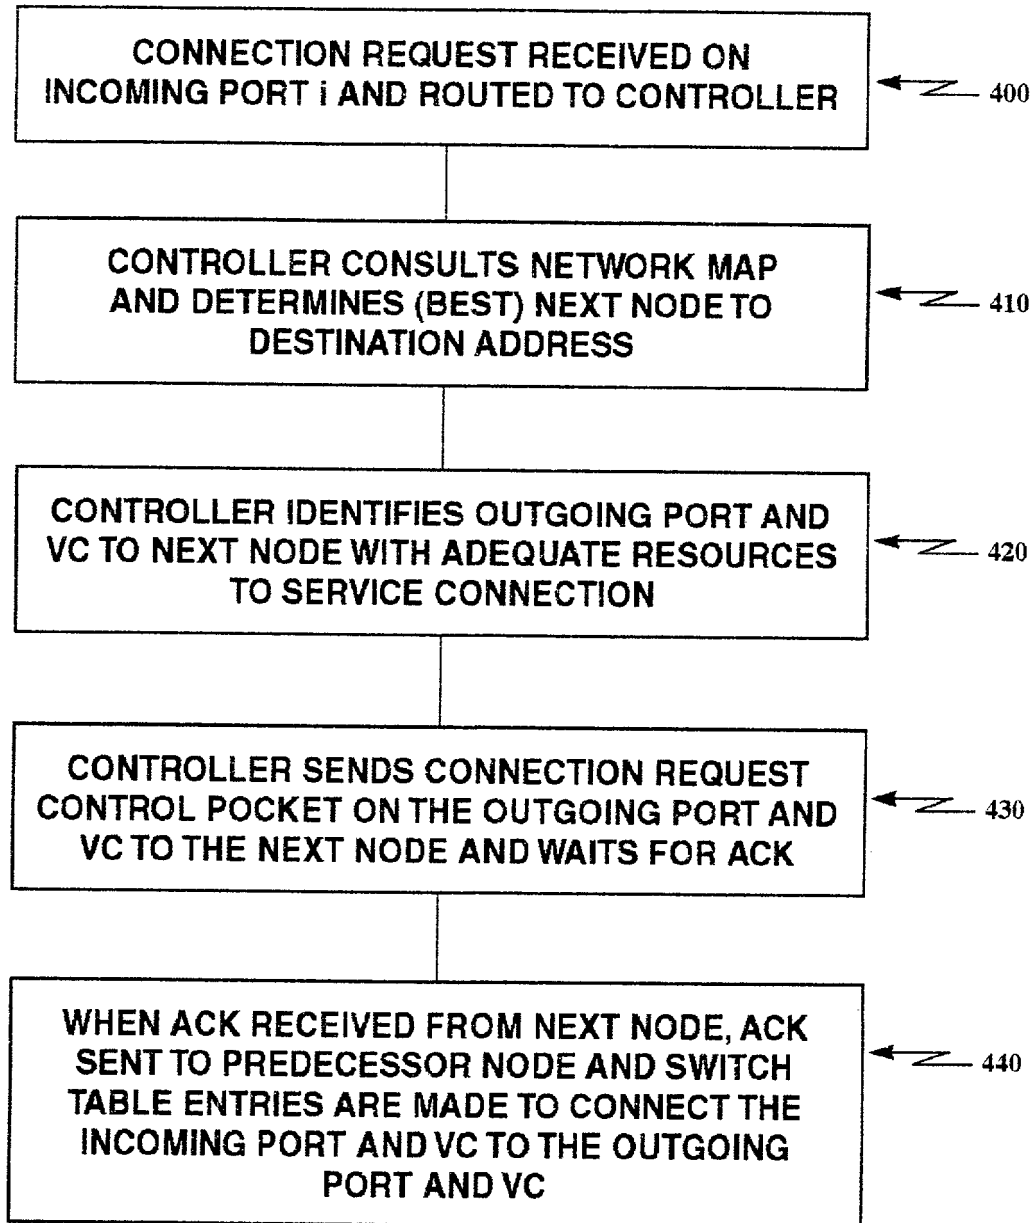


Figure 4

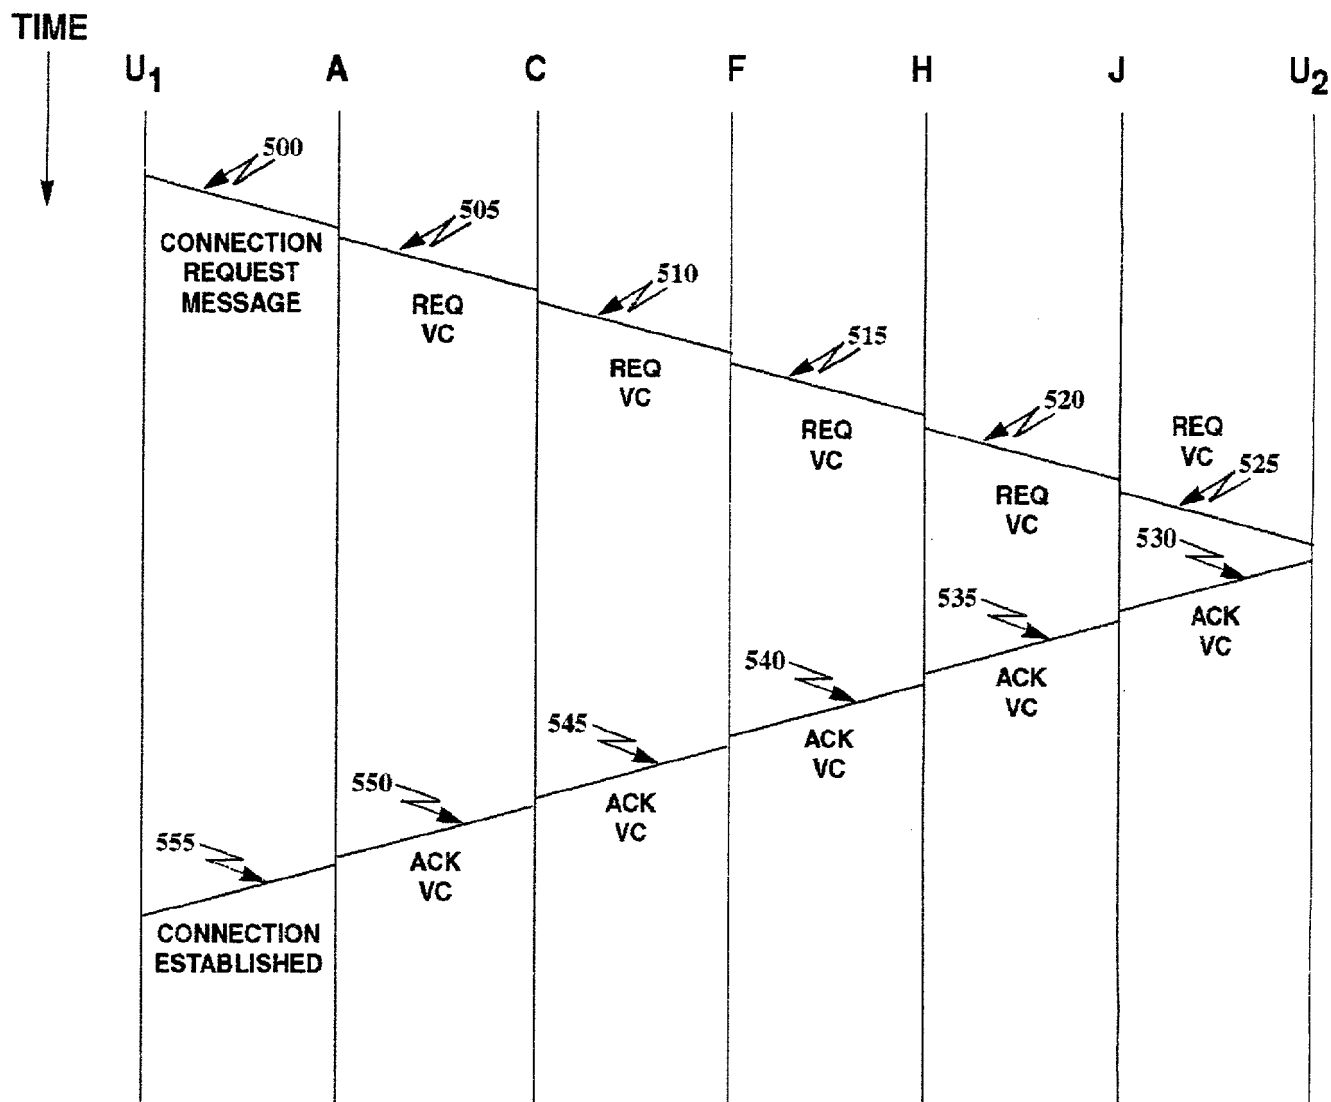


Figure 5

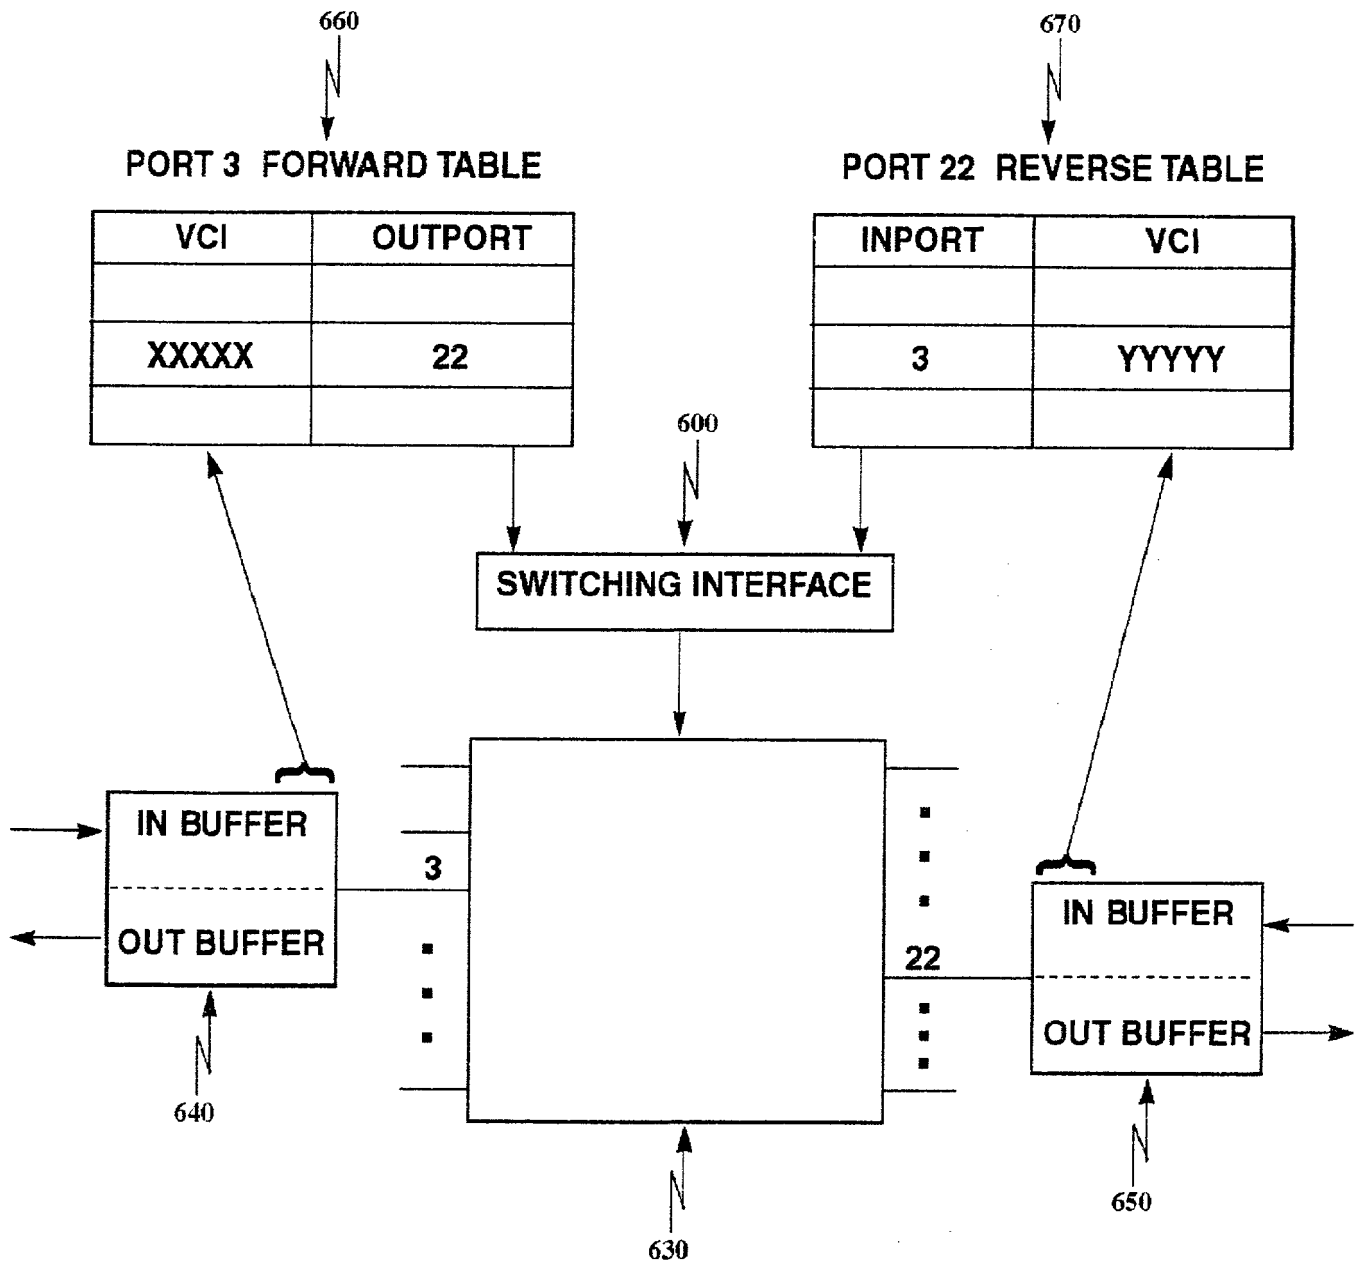


Figure 6

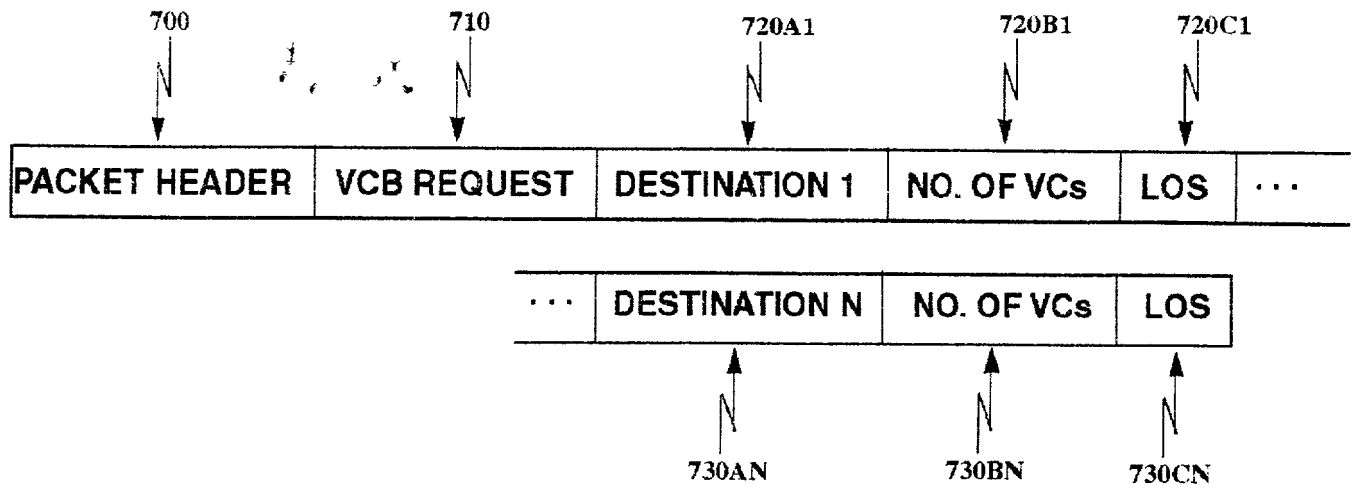


Figure 7A

PACKET HEADER	VCB REQUEST	SWITCH 2	32 VCs	LOS
---------------	-------------	----------	--------	-----

Figure 7B

PACKET HEADER	VCB REQUEST	SWITCH 4	8 VCs	LOS ₁
		SWITCH 5	8 VCs	LOS ₂
		SWITCH 3	9 VCs	LOS ₃

Figure 7C

464770" 2288880

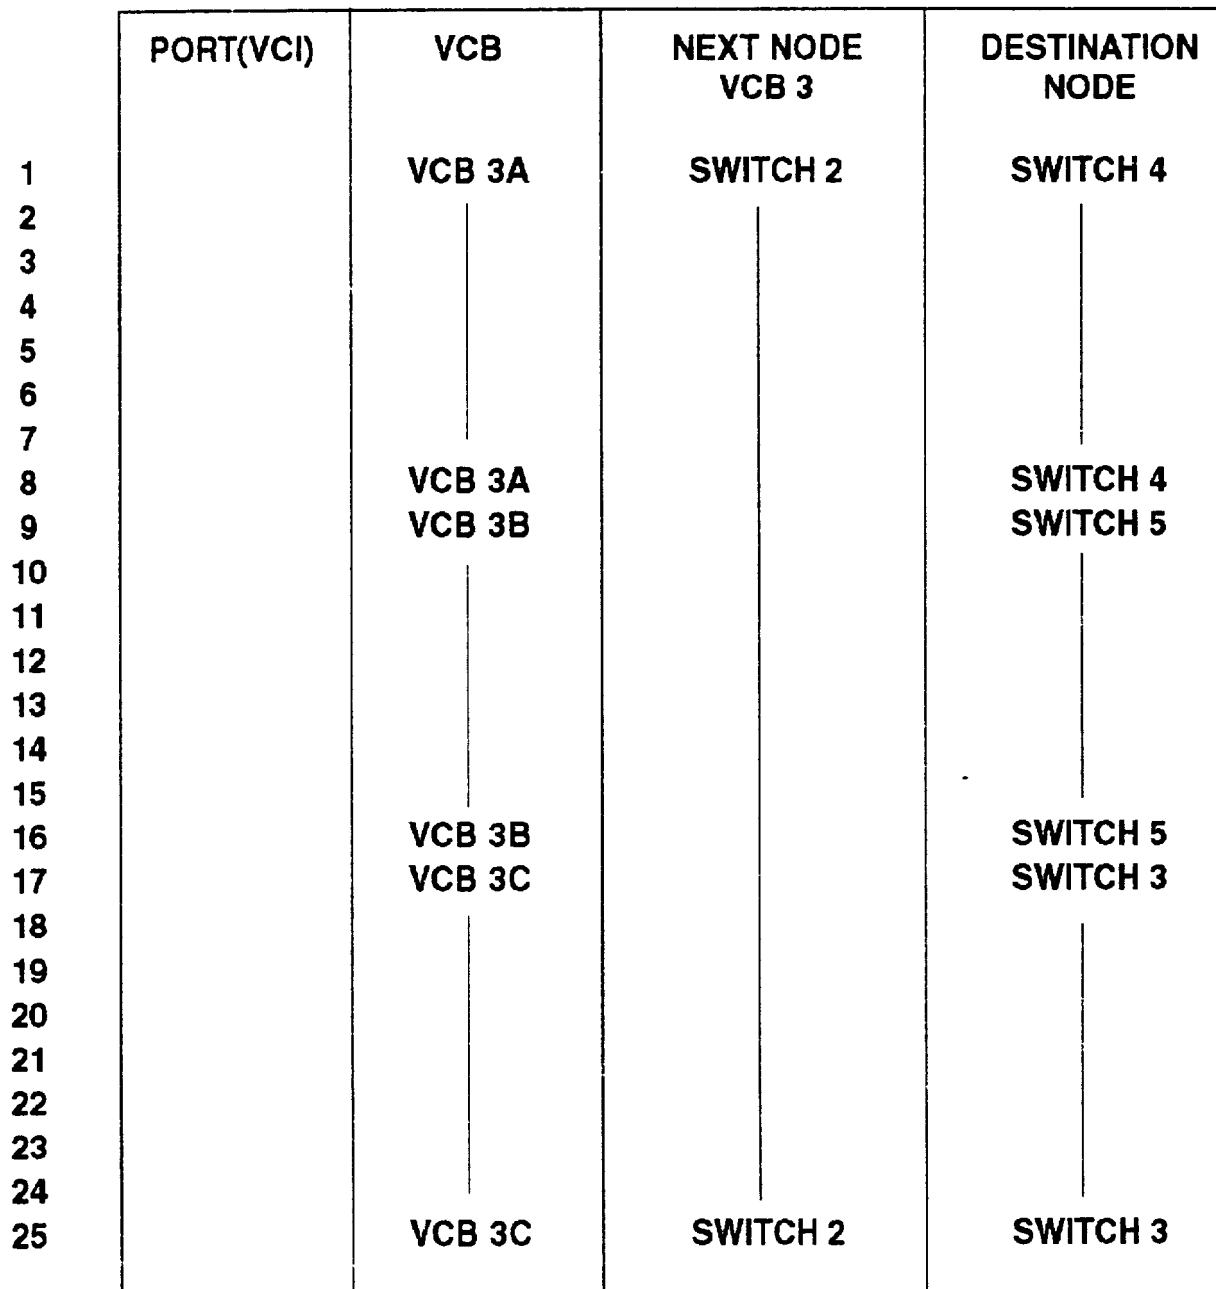


Figure 8A

PORT(VCI)	VCB	NEXT NODE VCB 3	DESTINATION NODE
3(102)	VCB 3A	SWITCH 2	SWITCH 4
3(103)			
3(104)			
3(105)			
3(106)			
3(111)			
3(112)			
3(113)	VCB 3A		SWITCH 4
5(103)	VCB 3B		SWITCH 5
5(107)			
5(109)			
5(112)			
5(113)			
5(114)			
5(115)			
5(116)	VCB 3B		SWITCH 5
2(101)	VCB 3C		SWITCH 3
2(102)			
2(103)			
2(104)			
2(105)			
6(107)			
6(108)			
6(109)			
6(110)	VCB 3C	SWITCH 2	SWITCH 3

Figure 8B



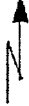


Figure 9

SWITCH 1 VCB DEFINITION TABLE

LEVEL 1	LEVEL 2	< PORT, VCI >
VCB 3	VCB 3A	PORT 5, 103-110; DESTINATION A
	VCB 3B	PORT 5, 112-114 PORT 6, 101-105; DESTINATION B
	VCB 3C	PORT 6, 107-115; DESTINATION C


1000


1010

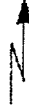

1020

Figure 10

SWITCH 1 CONTROLLER

FORWARD TABLE:

INPUT < PORT, VCI >	VCB	OUTPUT
3 (102-106, 111-113)	VCB 3A	ANY UNUSED
5 (103,107,109, 112-116)	VCB 3B	ANY UNUSED

REVERSE TABLE:

OUTPUT < PORT, VCI >	DESTINATION ADDRESS	VCB
3(102) 3(103) 3(104) 3(105) 3(106) 3(111) 3(112) 3(113) 5(103) 5(107) 5(109) 5(112) 5(113) . . . 6(110)	DESTINATION ADDRESS ↓	VCB 3 ↓

Figure 11

FORWARD TABLE:

INPUT < PORT, VCI >	VCB	OUTPUT
5 (103-111)	VCB 3A	ANY UNUSED (e.g. 1 of 8(109-116))
5 (112-114), 6(101-105)	VCB 3B	ANY UNUSED (e.g. 1 of 10(103-111))
6 (107-115)	VCB 3C	ANY UNUSED (e.g. 1 of 15(103-106), 17(101-105))

REVERSE TABLE:

OUTPUT < PORT, VCI >	VCB	INPUT < PORT, VCI >
ANY UNUSED	VCB 3	8(109-116), 10(103-111), 15(103-106), 17(101-105)

Figure 12

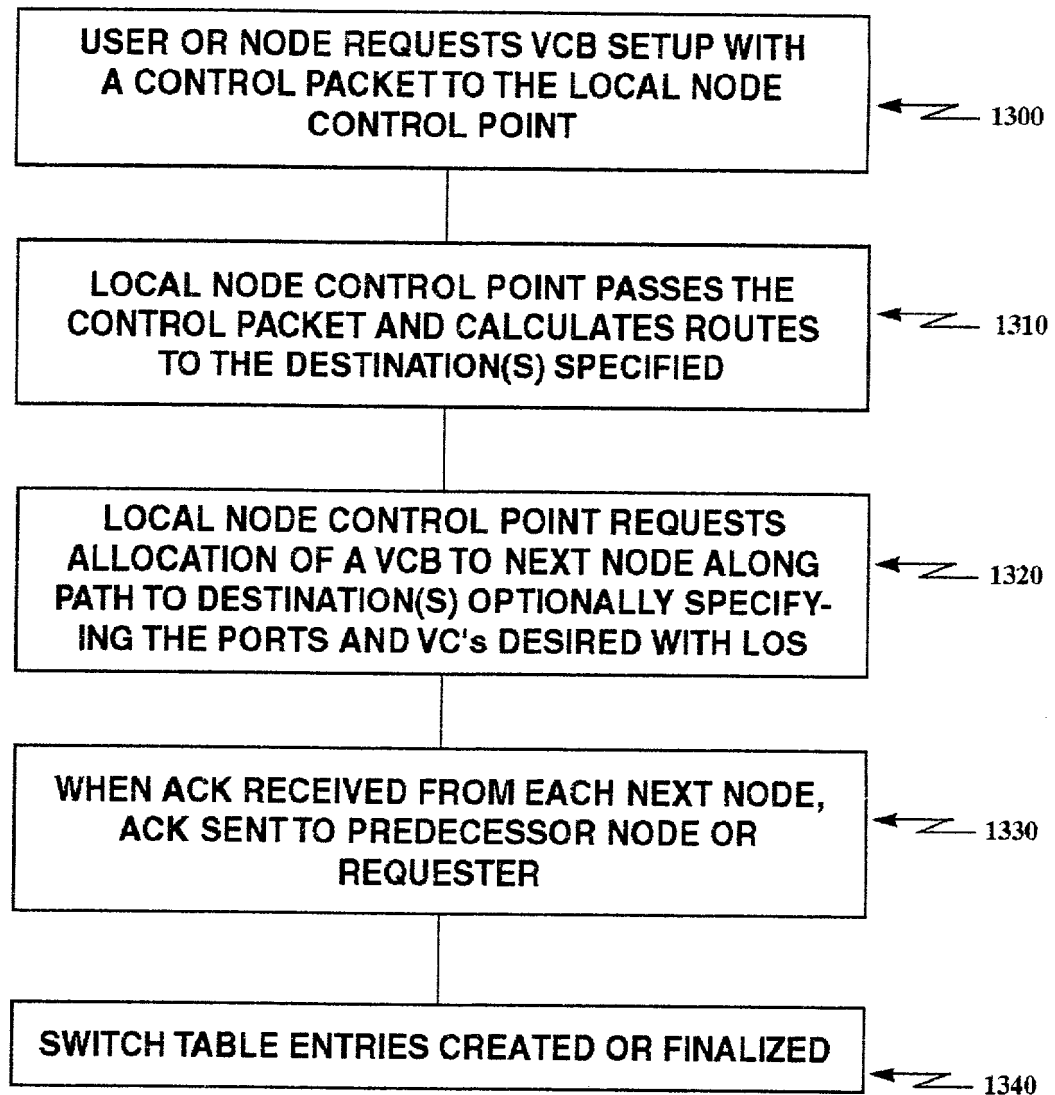


Figure 13

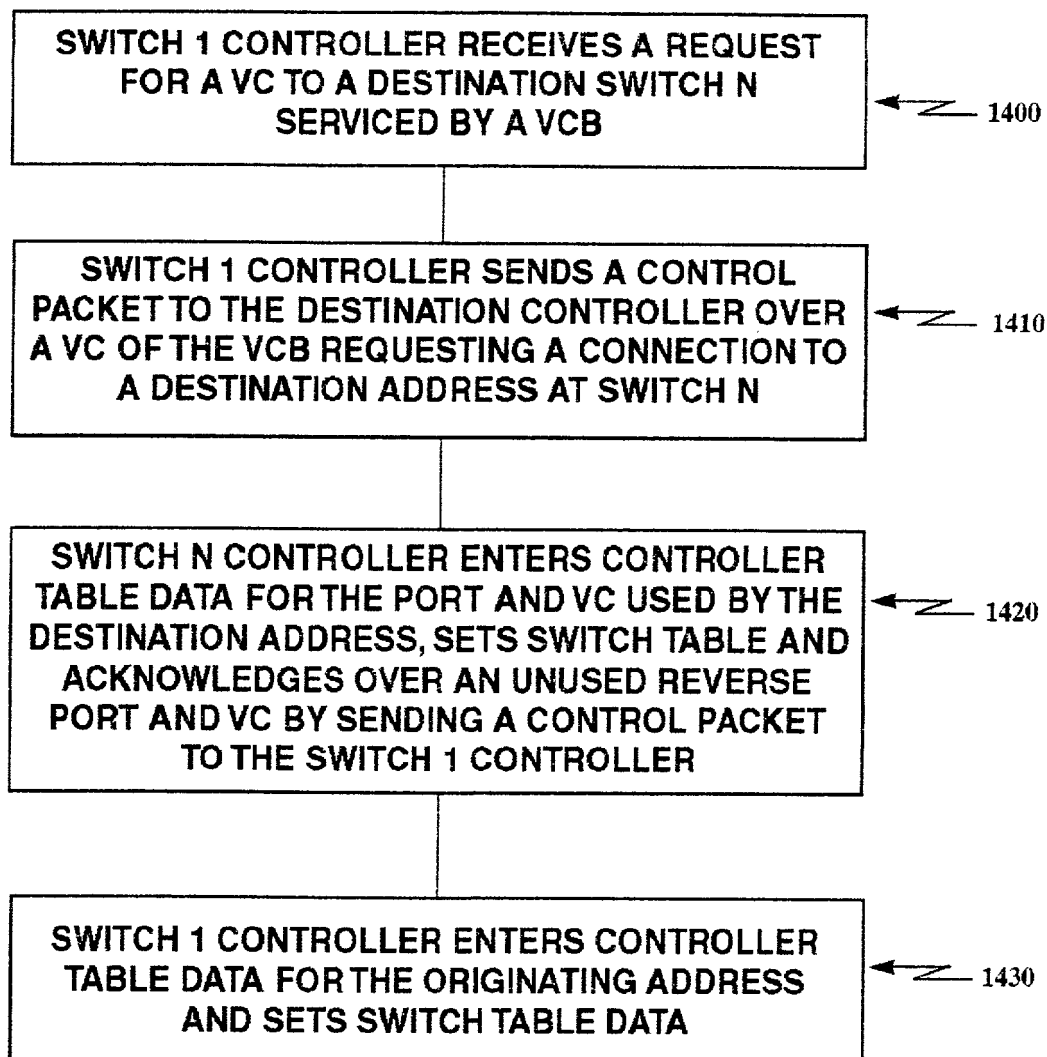


Figure 14

**VCB BREAKDOWN REQUEST RECEIVED ON
INCOMING PORT i AND ROUTED TO CONTROLLER**

← Z 1500

CONTROLLER CONSULTS TABLE FOR ENTRIES USING ALIAS OF VCB BEING BROKEN DOWN AND MARKS THEM FOR DELETION

← 1510

CONNECTION BREAKDOWN REQUEST FORWARDED TO NEXT NODE, IF ANY; OTHERWISE TABLE ENTRIES FOR VCB BEING BROKEN DOWN ARE DELETED AND ACK SENT TO REQUESTING NODE

← 1520

**WHEN ACK OF BREAKDOWN REQUEST IS
RECEIVED, TABLE ENTRIES ARE ACTUALLY
DELETED**

← Z 1530

Figure 15

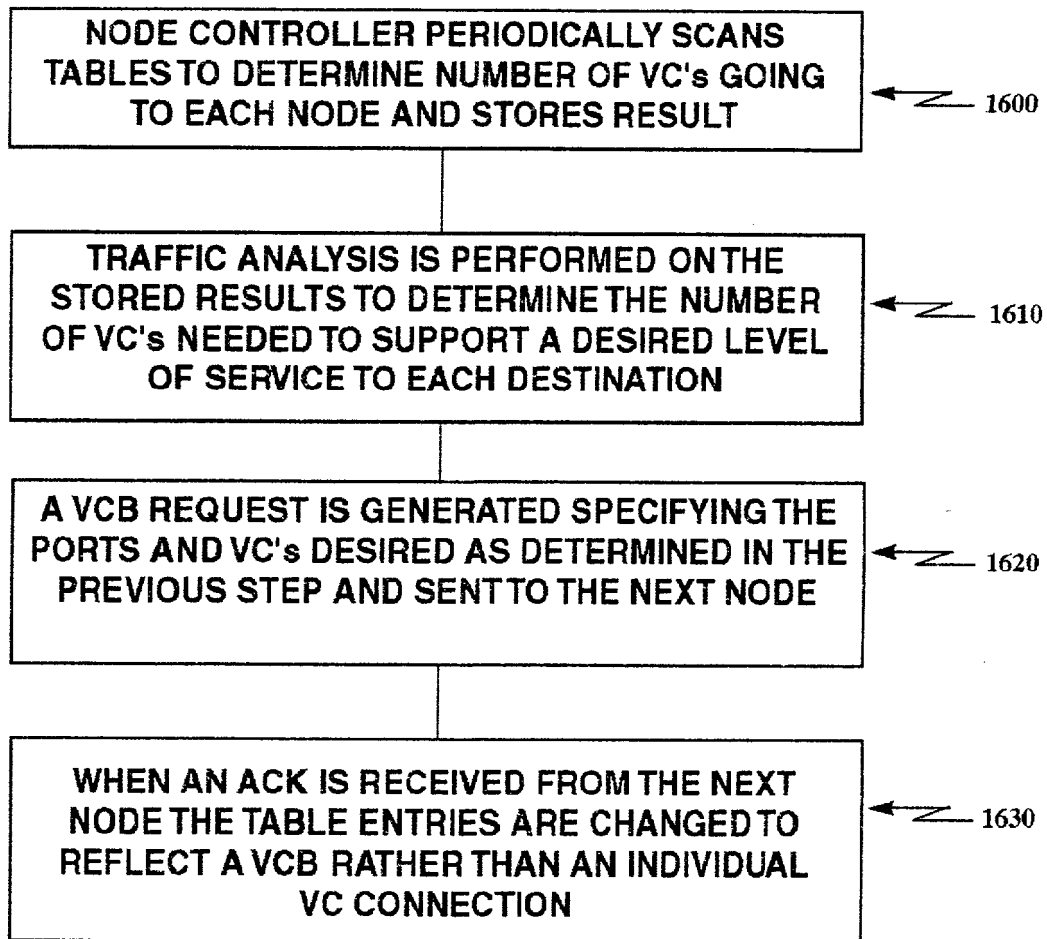


Figure 16




Figure 17

20250120 10:04:37

WHEN A FAST CONNECT (PACKET + REQUEST)
IS RECEIVED AT A NODE, A PORT AND A VC
FROM A VCB ARCH GOING TO THE DESTINATION
NODE IS SELECTED

← 1800

THE FAST CONNECT PACKET/REQUEST IS SENT
OVER THE SELECTED PORT AND VC TO THE
DESTINATION NODE

← 1810

THE FAST CONNECT PACKET/REQUEST IS ROUTED
TO THE DESTINATION NODE TO THE SPECIFIED
DESTINATION ADDRESS

← 1820

AN OPTIONAL FAST CONNECT ACK IS
RETURNED TO THE ORIGINATING ADDRESS IN
THE SAME WAY A FAST CONNECT
PACKET/REQUEST WOULD BE SENT

← 1830

Figure 18





Figure 19

PORT 1 SWITCH
TABLE

VCI _{IN}	PORT	VCI _{OUT}
⋮		
120	2	150
⋮		

← 2010



VCI _{IN}	PORT	VCI _{OUT}
⋮		
125	2	150
⋮		

← 2020

Figure 20

2010-02-26 14:00:00

Table 1. Demographic characteristics of the study population	
Age (years)	65.0 ± 1.5
Gender	Male 50.0%
Education (years)	12.0 ± 1.0
Marital status	Married 70.0%
Occupation	Retired 60.0%
Income (USD/month)	1,500.0 ± 200.0
Health status	Good 80.0%
Smoking status	Smoker 30.0%
Alcohol consumption	Occasional 40.0%
Exercise frequency	Regular 20.0%
Family size	3.0 ± 0.5
Living arrangement	Alone 10.0%
Religious affiliation	Christian 70.0%
Political affiliation	Democrat 60.0%
Travel frequency	Monthly 30.0%
Travel distance (miles)	50.0 ± 10.0
Travel mode	Car 70.0%
Travel companions	Alone 40.0%
Travel purpose	Leisure 60.0%
Travel satisfaction	Satisfied 80.0%
Travel frequency (times/month)	2.0 ± 0.5
Travel distance (miles/month)	100.0 ± 20.0
Travel mode (times/month)	1.0 ± 0.2
Travel companions (times/month)	0.5 ± 0.1
Travel purpose (times/month)	1.0 ± 0.2
Travel satisfaction (times/month)	1.0 ± 0.2
Travel frequency (times/year)	24.0 ± 6.0
Travel distance (miles/year)	1,200.0 ± 240.0
Travel mode (times/year)	12.0 ± 2.4
Travel companions (times/year)	6.0 ± 1.2
Travel purpose (times/year)	12.0 ± 2.4
Travel satisfaction (times/year)	12.0 ± 2.4




Figure 21A


FINAL STATE:

Figure 21B